IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

EX PARTE MARTIN CIESLAK

Application for Patent

Filed June 30, 2000

Application No. 09/608,802

FOR:

METHOD AND APPARATUS FOR REDIRECTING NETWORK TRAFFIC

APPEAL BRIEF

CERTIFICATE OF EFS-WEB TRANSMISSION

I hereby certify that this correspondence is being transmitted electronically through EFS-WEB to the Commissioner for Patents, P.O. Box 1450 Alexandria, VA 22313-1450 on August 27, 2007.

Signed: /Mia Mitchell-Haynes/

Mia Mitchell-Haynes

BEYER WEAVER LLP
Attorneys for Appellant

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST

The real party in interest is Cisco Technology, Inc., the assignee of the present application.

II. RELATED APPEALS AND INTERFERENCES

A Pre-Appeal Brief Request for Review was filed on April 30, 2007, along with the Notice of Appeal which precipitated the present appeal. A Pre-Appeal Brief conference was held. A copy of the Notice of Panel Decision from Pre-Appeal Brief Review dated June 25, 2007, is included herewith in the Related Proceedings Appendix.

III. STATUS OF THE CLAIMS

Claims 1-79 were filed with the present application. Claims 80-82 were subsequently added and claims 36 and 74 were cancelled in an amendment dated January 5, 2004. Claims 20-35, 37-39, 58-73, 75-77, 79, and 81-82 were cancelled in the amendment dated December 20, 2005. Claims 1-19, 40-57, 78, and 80 remain pending in this application.

Claims 1 and 80 stand rejected under 35 U.S.C. §101 because the claimed invention is asserted to be directed towards non-statutory subject matter. Claims 1-19, 40-57, 78 and 80 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Yates et al. (referred to herein as "Yates") in view of the document Cardellini et al., entitled "Dynamic Load Balancing on Web-Server System" (referred to herein as "Dynamic"), 1089-7801/99/IEEE, pages 28-39.

The rejection of each of claims 1-19, 40-57, 78 and 80 is appealed.

IV. STATUS OF AMENDMENTS

On January 30, 2007, the Examiner entered a final rejection. No amendments were filed in response to this final rejection.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

The present invention provides methods and apparatus which facilitate the redirection of data sent from a first processing device to a second processing device. Embodiments of the present invention may be utilized in a network such as illustrated below in Fig. 1:

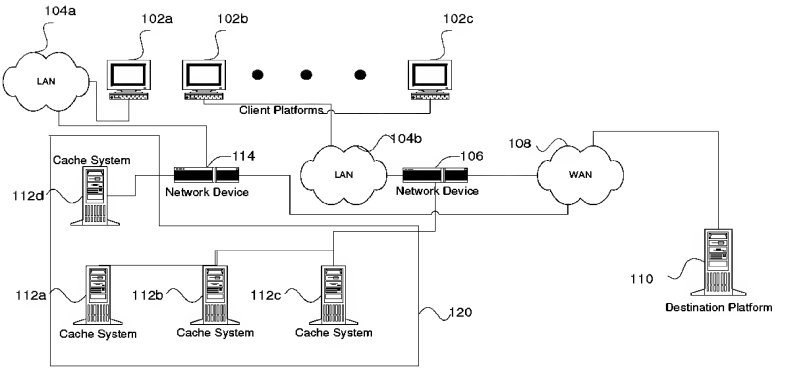

Fig. 1

As shown, a plurality of client machines 102 which are resident on local area networks (LAN) 104a and 104b communicate via router 106/router 114 and wide area network (WAN) 108, e.g., the internet, with server 110. In embodiments of the invention, this intermediary router 106 is configured to redirect certain traffic towards a selected cache system or cluster, such as cache systems 112a~112d, while the destination server 110 is being "spoofed." The router 106 selects a particular one of its associated cache systems based on traffic information received from each of its associated cache devices, *e.g.,* 112a~112d, and this received information specifies how to redirect received data based on at least a portion of such received data. For example, the traffic information may specify that a first set of destination IP addresses are assigned to cache system 112a; a second set of IP addresses to cache system 112b; a third set of IP addresses to cache system 112c; and a fourth set of IP addresses to cache system 112d. Thus, embodiments of the present invention allow a designated traffic handling system to specify how a router redirects data to selected ones of its associated *plurality* of traffic handling systems based on traffic information from each of the plurality of traffic handling systems and at least a portion of such received data, wherein the traffic information specifies how to redirect data based on at least a portion of the received data.

Thus, the router is receiving traffic information from each of its own associated cache systems that specifies how to redirect data based on some portion of such data, such as an IP address or port, and this traffic information from each of the router's cache systems is then used to determine how to redirect data based on the particular specified data portion. Accordingly, the router can then redirect data based on a information from each of its associated cache systems that specifies how to redirect based on a specified data portion. For example, this arrangement allows each cache system to flexibly adjust redirection of data to itself by its associated router based on the cache system's availability or processing abilities. For instance, a cache system may specify a narrow or a wider range of source IP addresses to be redirected to itself based on such cache system's load. Even though each cache system can indicate it's own preference for redirection in the traffic information that it sends to its associated router, the associated router can then redirect based on the total compilation of traffic information from each of its associated cache systems. That is, a router can then apportion data based on traffic information received

from each cache system that specifies redirection based on a particular data portion. Thus, when redirecting, the totality of the traffic information from a set of associated cache systems and how each cache system specified how to redirect based on a data portion can be considered.

Independent claim 1

Independent claim 1 recites a "method of facilitating redirection of data sent from a first processing device to a second processing device." See, for example, Fig. 2. The method includes "at a third processing device associated with a plurality of traffic handling systems, receiving traffic information from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system." See, for example, operation 202 of Fig. 2 and the corresponding description on page 12, lines 15-20, as well as page 7, lines 5-7. The method also includes "determining how to redirect data received by the third processing device to a selected traffic handling system based on the received traffic information from each of the associated traffic handling systems." See, for example, operations 204-208 in Fig. 2 and the corresponding description on page 13, lines 11-22, as well as page 6, lines 7-12. Claim 1 also recites "at the third processing device, receiving data from one or more first processing devices that are destined for one or more second processing devices." See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 22 through page 14, line 1, as well as page 6, line 24 through page 7, line 3. Claim 1 further recites "at the third processing device, redirecting the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data." See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 21 through page 14, line 1, as well as page 6, line 24 through page 7, line 3 and page 10, line 22 through page 11, line 1.

Independent claim 40

Independent claim 40 defines a "first computer system associated with a plurality of traffic handling systems and operable to facilitate redirection of data sent from a second computer system to a third computer system, the traffic being redirected to a selected traffic handling system." See network device 106 of Fig. 1 and page 10, lines 20-24, by way of examples. Claim 40 recites "a memory" and "a processor coupled to the memory." See 65 and 61 of Fig. 6 and page 22, lines 5-6, by way of examples. Claim 40 also recites "at least one of the memory and the processor are adapted to provide." See, for instance, page 22, line 21 through page 23, line 11. Claim 40 also recites "receiving traffic information from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system." See, for example, operations 202 of Fig. 2 and the corresponding description on page 12, lines 15-20, as well as page 7, lines 5-7. Claim 40 also recites that "communicating the received traffic information to at least a designated one of the associated traffic handling systems." See, for example, operation 204 in Fig. 2 and the corresponding description on page 13, lines 11-15. Claim 40 also recites "receiving traffic redirection information from the designated traffic handling system, the traffic redirection information specifying which data is to be redirected to which one of the plurality of traffic handling systems and being based on the communicated traffic information." See again, for example, operation 208 in Fig. 2 and the corresponding description on page 13, line 19 through

page 14, line 2. Claim 40 also recites "receiving data from one or more second computer systems that are destined for one or more third computer systems." See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 22 through page 14, line 1, as well as page 6, line 24 through page 7, line 3. Claim 40 also recites "redirecting the received data to selected one or more of the traffic handling systems so that the redirected data is apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data." See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 21 through page 14, line 1, as well as page 6, line 24 through page 7, line 3 and page 10, line 22 through page 11, line 1.

Independent claim 78

Independent claim 78 defines a "computer readable storage medium, that is in the form of magnetic media, optical media, or magneto-optical media, for facilitating redirection of data sent from a first processing device to a second processing device" and "comprising computer readable code that is stored on such computer readable storage medium and that is configured for performing the following operations." See, for example, page 24, lines 15-22. Claim 78 also recites "receiving traffic information, at a third processing device associated with a plurality of traffic handling systems, from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system." See, for example, operation 202 of Fig. 2 and the corresponding description on page 12, lines 15-20, as well as page 7, lines 5-7. Claim 78 also recites "communicating the traffic information to at least a designated one of the associated traffic handling systems." See, for example, operation 204 in Fig. 2 and the corresponding description on page 13, lines 11-15. Claim 78 also recites "receiving traffic redirection information, at the third processing device, from the designated traffic handling system, the traffic redirection information specifying which data is to be redirected to which one of the plurality of traffic handling systems and being based on the communicated traffic information." See again, for example, operation 208 in Fig. 2 and the corresponding description on page 13, line 19 through page 14, line 2. Claim 78 also recites "receiving data, at the third processing device, from one or more first processing devices that are destined for one or more second processing devices." See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 22 through page 14, line 1, as well as page 6, line 24 through page 7, line 3. Claim 78 also recites "redirecting, at the third processing device, the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data." See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 21 through page 14, line 1, as well as page 6, line 24 through page 7, line 3 and page 10, line 22 through page 11, line 1.

Independent claim 80

Independent claim 80 defines a apparatus for facilitating redirection of data sent from a first processing device to a second processing device, wherein the apparatus is associated with a plurality of traffic handling systems. See 65 and 61 of Fig. 6 and page 22, lines 5-6, by way of examples. Claim 80 recites "means for receiving traffic information from each of the associated

traffic handling systems, wherein the traffic information received from each of the associated traffic handling systems specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system." See 65 and 61 of Fig. 6 and page 22, lines 5-6, by way of examples. Additionally, see, for example, operation 202 of Fig. 2 and the corresponding description on page 12, lines 15-20, as well as page 7, lines 5-7. Claim 80 also recites that "means for determining how to redirect data received by the apparatus to a selected traffic handling system based on the received traffic information." See 65 and 61 of Fig. 6 and page 22, lines 5-6 and operations 204-208 in Fig. 2 and the corresponding description on page 13, lines 11-22, as well as page 6, lines 7-12, by way of examples. Claim 80 also recites "means for receiving data, at the third processing device, from one or more first processing devices that are destined for one or more second processing devices." See 65 and 61 of Fig. 6 and page 22, lines 5-6, by way of examples. See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 22 through page 14, line 1, as well as page 6, line 24 through page 7, line 3. Claim 80 also recites "means for redirecting, at the third processing device, the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data." See 65 and 61 of Fig. 6 and page 22, lines 5-6, by way of examples. See again, for example, operation 210 in Fig. 2 and the corresponding description on page 13, line 21 through page 14, line 1, as well as page 6, line 24 through page 7, line 3 and page 10, line 22 through page 11, line 1.

Dependent claim 2

Dependent claim 2 recites "wherein the determination of redirecting data is accomplished by communicating the traffic information to at least a designated one of the associated traffic handling systems." See, for example, operation 204 in Fig. 2 and the corresponding description on page 13, lines 11-15. Claim 2 also recites "and at the third processing device, receiving traffic redirection information from the designated traffic handling system, the traffic redirection information being based on the communicated traffic information." See again, for example, operation 208 in Fig. 2 and the corresponding description on page 13, line 19 through page 14, line 2.

Dependent claim 9

Dependent claim 9 further recites "wherein the fields are selected from a group consisting of a source IP field, a destination IP field a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field." See 304 of Fig. 3 and page 14, line 15 through page 15, line 5, among other places.

Dependent claim 10

Dependent claim 10 further recites "wherein each field indicates that a corresponding field of a packet received in the third processing device is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being

associated with the selected traffic handling system." See page 14, lines 18-20 and page 15, lines 1-4, lines 1-4, among other places.

Dependent claim 11

Dependent claim 11 further recites "wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems." See page 14, lines 18-20 and page 15, lines 1-4, among other places.

Dependent claim 12

Dependent claim 12 further recites "wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the third processing device." See 304 of Fig. 3 and page 14, line 15 through page 15, line 5, among other places.

Dependent claim 13

Dependent claim 13 further recites "wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field." See 304 of Fig. 3 and page 14, line 15 through page 15, line 5.

Dependent claim 47

Dependent claim 47 further recites "wherein the fields are selected from a group consisting of a source IP field, a destination IP field, a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field." See 304 of Fig. 3 and page 14, line 15 through page 15, line 5, among other places.

Dependent claim 48

Dependent claim 48 further recites "wherein each field indicates that a corresponding field of a packet received in the first computer system is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system." See page 14, lines 18-20 and page 15, lines 1-4, among other places.

Dependent claim 49

Dependent claim 49 further recites "wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems." See page 14, lines 18-20 and page 15, lines 1-4, among other places.

Dependent claim 50

Dependent claim 50 further recites "wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the computer system." See 304 of Fig. 3 and page 14, line 15 through page 15, line 5, among other places.

Dependent claim 51

Dependent claim 51 further recites "wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field." See 304 of Fig. 3 and page 14, line 15 through page 15, line 5, among other places.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

(Ground 1): Whether claims 1 and 80 are patentable under 35 U.S.C. §101 so that the claimed invention is directed towards statutory subject matter.

(Ground 2): Whether claims 1-19, 40-57, 78 and 80 are patentable under 35 U.S.C. §103(a) over Yates et al. (referred to herein as "Yates") in view of the document Cardellini et al., entitled "Dynamic Load Balancing on Web-Server System" (referred to herein as "Dynamic"), 1089-7801/99/IEEE, pages 28-39.

VII. ARGUMENT

(Ground 1):

(1) Independent claims 1 and 80

The Examiner has rejected claims 1 and 80 under 35 U.S.C. §101 because the claimed invention is asserted to be directed towards non-statutory subject matter. The Examiner asserts that these claims are directed towards "software, which is not implemented on any type of hardware and fails to produce a tangible result like saving the data to memory." Claims 1 and 80 include techniques or mechanisms for "redirecting the received data to a selected one or more of the traffic handling systems" based on determined factors. It is respectfully submitted that the redirection of data to a selected traffic handling system is a tangible action that accomplishes the tangible result of redirecting data to a tangible device, *i.e.,* a selected traffic handling system. Accordingly, it is submitted that claims 1 and 80 are directed towards patentable subject matter.

(Ground 2):

The Examiner rejected claims 1-19, 40-57, 78 and 80 under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 6,167, 438 to Yates et al. (referred to herein as "Yates") in view of Dynamic Load Balancing on Web-Server System (Dynamic). Applicants traverse this rejection with respect to different claim groups as follows.

(1) Independent claims 1 and 80

Claim 1 is directed towards a "method of facilitating redirection of data sent from a first processing device to a second processing device." Claim 1 also recites "at a third processing device associated with a plurality of traffic handling systems, receiving traffic information from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system." Thus, traffic information is received by an intermediary device from each of the intermediary device's associated traffic handling systems, e.g., cache systems. This traffic information that is received from each cache system specifies which received data, based on at least a portion of such received data, is to be redirected to the each cache system. Claim 1 further requires "determining how to redirect data received by the third processing device to a selected traffic handling system based on the received traffic information from each of the associated traffic handling systems." In other words, determining how to redirect data to a selected cache system is based on the traffic information received from each of the cache systems.

Claim 1 also requires "at the third processing device, receiving data from one or more first processing devices that are destined for one or more second processing devices" and "at the third processing device, redirecting the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data." Claim 80 is directed towards mechanisms for performing the operations of claim 1. Thus, data that is received by the intermediary device is redirected from its destination to the selected cache system, wherein determining how to redirect to such selected cache system is based on traffic information, from <u>each</u> of the associated cache systems, that indicates how to redirect based on a data portion, e.g., a particular cache system specifies in its traffic information that particular IP addresses are to be redirected to the particular cache system and the other cache systems specify other IP addresses (distinct or overlapping) for redirection to themselves. That is, redirection in the present invention relies on information from each of its associated cache systems that specifies how redirection is accomplished to such each cache system based on a received data portion, such as a particular IP address range. These various pieces of redirection information from each of the cache systems can then be used to determine how to apportion data to the associated cache systems in an intelligent manner since redirection is based on information from each of the cache systems and such information <u>specifies how to redirect</u> to a selected cache system <u>based on a data portion</u>. That is, redirection is based on each cache system specifying how to redirect data to such each cache system based on a data portion so that each cache system can tailor its redirection specification to its own individual needs, i.e., availability or load, by simply indicating redirection being based on a particular data portion, e.g., such as a particular range of addresses.

Although the cited reference Yates is directed towards a distributed cache system, Yates fails to teach or suggest a single device that determines how to redirect data that is received by such single device based on traffic information that specifies how to redirect data based on at least a portion of such received data, such as the destination field, wherein such traffic information is received by each of a plurality of traffic handling devices that are associated with such single device. Specifically, Yates teaches an arrangement of routers that are distributed in a path that a request would take as such request is sent from a client to a home server. See Fig. 1 and Col. 7, Lines 18-24. Each router has at most a single cache. In Fig. 1, see each individual cache server 16 of some of the routers 14, such as cache 16-8 associated with router 14-8 and cache 16-4 associated with router 14-4. Yates goes on to teach that each router either intercepts data being sent between the client and the home server such that the intercepted data is handled by the specific router's single associated cache server or forwards the data towards the home server. See col. 7, lines 52-56. The router's decision to intercept its received data is based on filter code that is updated by the router's single cache. See col. 7, lines 46-56. This filter code merely indicates whether the router's single cache has the requested document. See col. 9, lines 1-7: "If the filter 26 indicates that the requested document is cached and can be serviced locally, then the packet is intercepted and passed to the resource manager 24. Otherwise, the packet is passed on to the next hop towards the destination home server 20."

Thus, a router in the path of a packet merely decides whether to redirect a received packet to its single cache system based on a filter code which merely indicates whether such associated cache has the requested document. Yates fails to teach or suggest any other information that can be assessed from such filter code, besides an indication as to whether a single associated cache contains the requested document, so that such assessment can aid in redirection determination by such router. Although Yates does generally describes load balancing techniques that are implemented by the cache servers to distribute documents among themselves based on information shared between the cache servers in section 4 on column 13 through 16. Yates fails to teach a technique (or mechanism) for apportioning *redirected* data among such cache servers based on traffic information received from each of the cache servers, wherein the traffic information specifies how to redirect based on a received data portion, in the manner claimed. Instead, it would appear that filter codes are merely used to indicate whether a router's cache system has a requested document after such load balancing of documents has been accomplished by the cache systems. That is, the filter codes do not specify how to redirect to a selected server based on a received data portion, in the manner claimed. In sum, Yates fails to teach or suggest that data is redirected to a cache system that is selected from a plurality of associated cache systems, wherein such redirection is based on traffic information that is received from each of the associated cache systems and that specifies how to redirect data to such each cache system based on at least a portion of such received data. In contrast, the filter code of each router in Yates simply indicates whether the single associated cache contains the requested document.

The Examiner admits that Yates fails to teach "based on the traffic information from each associated traffic handling system and at least a portion of the received data." The Examiner asserts that the secondary reference Dynamic teaches this limitation. Although Dynamic appears to teach redirecting data to a plurality of servers based on load information from such servers (e.g., Page 34), it is respectfully submitted that Dynamic fails to teach or suggest receiving traffic information from each of a plurality of cache servers that specifies how to redirect received data to each cache server based on at least a portion of such received data, in the manner claimed. Nor does Dynamic teach redirecting data based on information from each of a plurality of servers that specifies how to redirect received data to such each server based on at least a portion of such

received data, as well as based on at least a portion of such received data, in the manner claimed. In contrast, Dynamic teaches "[s]ince the dispatcher must be aware of the server load, each server periodically reports the number of received requests per second...The dispatcher then selects the least-loaded server." See last Paragraph on Page 34 (Emphasis Added). That is, redirection as taught by Dynamics is simply based on the load reported by a server, rather than on information that "specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system", in the manner claimed.

In more detail, Dynamic teaches two redirection techniques. The first technique is a *static* technique that uses a hash of a portion of the data to redirect the data to specific servers. See page 36, first column, lines 26-32. That is, this redirection scheme does not depend on information received from each of a plurality of associated cache servers, in the manner claimed. The second technique uses load information from the associated cache servers to determine redirection so as to balance the data load among the cache servers. Supra, lines 33-36. However, this second technique does not account for information from each server that specifies how to redirect data to such each server based on a portion of the received data, in the manner claimed. Although Dynamic may be argued to teach using load information from a plurality of cache servers to determine redirection, Dynamic cannot be said to teach using received information from a plurality of cache servers, wherein the received information from each server specifies how to redirect the received data to such each server based on a portion of said received data, in the manner claimed. In sum, the cited references both fail to teach or suggest mechanisms or techniques for redirecting data to selected ones of a plurality of traffic handling systems based on traffic information from each of the plurality of traffic handling systems and at least a portion of such received data, wherein the traffic information (received from the traffic handling systems) specifies how to redirect data to such each traffic handling system **based on at least a portion of the received data,** in the manner claimed.

Accordingly, it is respectfully submitted that the claims 1 and 80 are patentable over the cited art.

(2) Independent claims 40 and 78 and dependent claim 2

Claims 40 and 78 recite ""receiving traffic information, at a third processing device associated with a plurality of traffic handling systems, from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system." Claims 40 and 78 also recite "communicating the traffic information to at least a designated one of the associated traffic handling systems" and "receiving traffic redirection information, at the third processing device, from the designated traffic handling system, the traffic redirection information specifying which data is to be redirected to which one of the plurality of traffic handling systems and being based on the communicated traffic information." Claims 40 and 78 further recite "redirecting, at the third processing device, the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data." Dependent claim 2 also recites such limitations.

Yates fails to teach or suggest determining redirection by (i) communicating traffic information to a designated cache system, wherein such traffic information was received by an intermediary router as described with respect to claim 1, and (ii) receiving redirection information from such designated cache system specifying which data is to be redirected to which one of the cache systems and being based on such traffic information, which specified how to redirect based on a received data portion, in the manner claimed. Since Yates merely redirects based on whether a filter code indicates whether a single associated cache server has the requested document, Yates necessarily fails to teach or suggest such a specific mechanism for determining redirection by sending traffic information, that was received from each of a plurality of associated cache servers and that specifies how to redirect to each of these cache systems based on a received data portion, to a designated cache system, as claimed.

Thus, claims 40, 78, and 2 are further patentable over the cited art.

(3) Dependent claims 9 and 47

Dependent claims 9 and 47 further recites "wherein the fields are selected from a group consisting of a source IP field, a destination IP field a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field." The intervening claims specify that these one or more fields of the received data are used to specify which data is to be redirected to the selected cache system. That is, redirection is based on traffic information from each cache system that specifies which fields of the received data to use for redirection to such each cache system. Since Yates teaches redirection of received data based on filter codes that merely specify whether the single associated cache system contains the requested document, Yates necessarily fails to teach or suggest basing redirection on information that is received from each of the associated cache systems and that specifies which fields (as listed in claims 9 and 47) of the received data to use, in the manner claimed. Although Yates does generally describes load balancing techniques that are implemented by the cache servers to distribute documents among themselves based on information shared between the cache servers in section 4 of column 13 through 16, Yates fails to teach a technique (or mechanism) for *redirecting* data among such cache servers based on traffic information received from each of the cache servers, wherein the traffic information specifies how to redirect based on use of one or more of the claimed fields of the received data, in the manner claimed. Instead, it would appear that filter codes are merely used to indicate whether a router's cache system has a requested document after such load balancing of documents has been accomplished by the cache systems. That is, the filter codes do not specify how to redirect to a selected server based on the use of one or more claimed fields of the received data in the manner claimed. Although the secondary reference Dynamics appears to teach redirection by a static hashing of certain fields of received data, such hashing is not specified by each of a plurality of associated cache systems, in the manner claimed. Thus, claims 9 and 47 are further patentable over the cited references.

(4) Dependent claims 10 and 48

Dependent claims 10 and 48 further recites "wherein each field indicates that a corresponding field of a packet received in the third processing device is to be used to generate an index to a

table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system." Although Yates does generally describes load balancing techniques that are implemented by the cache servers to distribute documents among themselves based on information shared between the cache servers in section 4 of column 13 through 16, Yates fails to teach a technique (or mechanism) for *redirecting* data among such cache servers based on traffic information received from each of the cache servers, wherein the traffic information specifies how to redirect based on use of one or more of the claimed fields of the received data so as to generate an index from such field, in the manner claimed. Instead, it would appear that filter codes are merely used to indicate whether a router's cache system has a requested document after such load balancing of documents has been accomplished by the cache systems. That is, the filter codes do not specify how to redirect to a selected server based on the use of one or more claimed fields of the received data so as to generate an index from such field in the manner claimed. Although the secondary reference Dynamics appears to teach redirection by a static hashing of certain fields of received data, such hashing is not specified by each of a plurality of associated cache systems, in the manner claimed. Thus, claims 10 and 48 are further patentable over the cited references.

(5) Dependent claims 11 and 49

Dependent claims 11 and 49 nfurther recites "wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems." Although Yates does generally describes load balancing techniques that are implemented by the cache servers to distribute documents among themselves based on information shared between the cache servers in section 4 of column 13 through 16, Yates fails to teach a technique (or mechanism) for *redirecting* data among such cache servers based on traffic information received from each of the cache servers, wherein the traffic information specifies how to redirect based on use of one or more of the claimed fields of the received data so as to generate an index from a hashed value of such field, in the manner claimed. Instead, it would appear that filter codes are merely used to indicate whether a router's cache system has a requested document after such load balancing of documents has been accomplished by the cache systems. That is, the filter codes do not specify how to redirect to a selected server based on the use of one or more claimed fields of the received data so as to generate an index from a hash value of such field in the manner claimed. Although the secondary reference Dynamics appears to teach redirection by a static hashing of certain fields of received data, such hashing is not specified by each of a plurality of associated cache systems, in the manner claimed. Thus, claims 11 and 49 are further patentable over the cited references.

(6) Dependent claims 12 and 50

Dependent claims 12 and 50 further recites "wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the third processing device." The intervening claims specify that these one or more fields of the received data are used to specify which data is to be redirected to the selected cache system. That is, redirection is based on traffic information from each cache system that specifies port identifiers of the received data to use for redirection to such each cache system. Since Yates merely determine redirection of received data

based on filter codes that merely specify whether the single associated cache system contains the requested document, Yates necessarily fails to teach or suggest basing redirection on information from each of the cache system that specifies port identifiers (as listed in claims 12 and 50) of the received data to use, in the manner claimed. Although Yates does generally describes load balancing techniques that are implemented by the cache servers to distribute documents among themselves based on information shared between the cache servers in section 4 of column 13 through 16, Yates fails to teach a technique (or mechanism) for *redirecting* data among such cache servers based on traffic information received from each of the cache servers, wherein the traffic information specifies how to redirect based on use of port identifiers of the received data, in the manner claimed. Instead, it would appear that filter codes are merely used to indicate whether a router's cache system has a requested document after such load balancing of documents has been accomplished by the cache systems. That is, the filter codes do not specify how to redirect to a selected server based on the use of port identifiers of the received data in the manner claimed. Although the secondary reference Dynamics appears to teach redirection by a static hashing of certain fields of received data, such hashing is not specified by each of a plurality of associated cache systems, in the manner claimed. Thus, claims 12 and 50 are further patentable over the cited references.

(7) Dependent claims 13 and 51

Dependent claims 13 and 51 further recites "wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field." The filter codes of Yates that are used for redirection do not specify how to redirect to a selected server based on the use of port identifiers (as listed in claims 13 and 51) of the received data in the manner claimed. Although the secondary reference Dynamics appears to teach redirection by a static hashing of certain fields of received data, such hashing is not specified by each of a plurality of associated cache systems, in the manner claimed. Thus, claims 13 and 51 are further patentable over the cited references.

Conclusion

In view of the foregoing, it is respectfully submitted that the rejections of claims 1-19, 40-57, 78, and 80 of the present application under 35 U.S.C. §101 and 35 U.S.C. §103(a) should be reversed.

Respectfully Submitted,
BEYER WEAVER LLP

/Mary R. Olynick/

Mary R. Olynick
Registration No. 42,163

VIII. CLAIMS APPENDIX

CLAIMS ON APPEAL

1. (previously presented) A method of facilitating redirection of data sent from a first processing device to a second processing device, the method comprising:

at a third processing device associated with a plurality of traffic handling systems, receiving traffic information from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system;

determining how to redirect data received by the third processing device to a selected traffic handling system based on the received traffic information from each of the associated traffic handling systems;

at the third processing device, receiving data from one or more first processing devices that are destined for one or more second processing devices; and

at the third processing device, redirecting the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data.

2. (Previously presented) A method as recited in claim 1, wherein the determination of redirecting data is accomplished by:

communicating the traffic information to at least a designated one of the associated traffic handling systems; and

at the third processing device, receiving traffic redirection information from the designated traffic handling system, the traffic redirection information being based on the communicated traffic information.

3. (Previously Presented) A method as recited in claim 2, further comprising:

at the third processing device, building or updating a data structure based on the received traffic information, wherein the traffic information is communicated to the designated traffic handling system within the data structure.

4. (previously presented) A method as recited in claim 2, further comprising:

in the third processing device, receiving a packet from a first processing device destined for a second processing device; and

redirecting the packet to a selected one of the traffic handling systems based on the traffic redirection information.

5. (previously presented) A method as recited in claim 4, further comprising:

receiving the packet back in the third processing device after redirecting it to the selected traffic handling system;

determining that the packet is to be sent to the packet's original destination address instead of being redirected to the selected traffic handling system; and

sending the packet from the third processing device to its original destination.

6. (Original) A method as recited in claim 5, wherein it is determined that the packet is to be sent to the packet's original destination by determining that the packet is encapsulated and de-encapsulating the packet prior to sending the packet to its original destination.

7. (Previously presented) A method as recited in claim 1, wherein the traffic information sent from a selected traffic handling system to the third processing device includes service options specifying which data is to be redirected to the selected traffic handling system.

8. (Original) A method as recited in claim 7, wherein the service options include a plurality of fields that are configurable to indicate that one or more fields of a packet received in the third processing device are to be used to determined redirection of packets to the selected traffic handling system.

9. (Original) A method as recited in claim 8, wherein the fields are selected from a group consisting of a source IP field, a destination IP field a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field.

10. (Original) A method as recited in claim 9, wherein each field indicates that a corresponding field of a packet received in the third processing device is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system.

11. (Original) A method as recited in claim 10, wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems.

12. (Original) A method as recited in claim 8, wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the third processing device.

13. (Original) A method as recited in claim 8, wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field.

14. (Original) A method as recited in claim 13, wherein the fields includes a source/destination field to indicate whether the port identifiers of the received traffic are source ports or destination ports.

15. (Original) A method as recited in claim 6, further comprising:

in the third processing device, receiving a packet from the first processing device destined for the second processing device; and

when one or more port identifiers of the received packet matches a corresponding set field of the service options of the selected traffic handling system, redirecting the packet to the selected traffic handling system.

16. (Original) A method as recited in claim 1, wherein the traffic information sent from a selected traffic handling system to the third processing device includes security options for specifying an authentication level for messages communicated between the third processing device and the selected traffic handling system.

17. (Original) A method as recited in claim 16, wherein the security options are configurable to select no authentication for messages communicated between the third processing device and the selected traffic handling system.

18. (Original) A method as recited in claim 16, wherein the security options are configurable to require a predetermined password encoded within messages communicated between the third processing device and the selected traffic handling system.

19. (Original) A method as recited in claim 1, wherein the traffic information sent from a selected traffic handling system includes identifying information for the selected traffic handling system.

20-39. (Cancelled)

40. (Previously presented) A first computer system associated with a plurality of traffic handling systems and operable to facilitate redirection of data sent from a second computer system to a third computer system, the traffic being redirected to a selected traffic handling system, the computer system comprising:

a memory; and

a processor coupled to the memory,

wherein at least one of the memory and the processor are adapted to provide:

receiving traffic information from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system;

communicating the received traffic information to at least a designated one of the associated traffic handling systems;

receiving traffic redirection information from the designated traffic handling system, the traffic redirection information specifying which data is to be redirected to which one of the plurality of traffic handling systems and being based on the communicated traffic information;

receiving data from one or more second computer systems that are destined for one or more third computer systems; and

redirecting the received data to selected one or more of the traffic handling systems so that the redirected data is apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data.

41. (Original) A computer system as recited in claim 40, wherein at least one of the memory and the processor are further adapted to provide:

building or updating a data structure based on the received traffic information, wherein the traffic information is communicated to the designated traffic handling system within the data structure.

42. (Original) A computer system as recited in claim 40, wherein at least one of the memory and the processor are further adapted to provide:

receiving a packet from a second computer system destined for a third computer system; and

redirecting the packet to a selected one of the traffic handling systems based on the traffic redirection information.

43. (Previously presented) A computer system as recited in claim 42, wherein at least one of the memory and the processor are further adapted to provide:

receiving the packet back in the first computer system after redirecting it to the selected traffic handling system;

determining that the packet is to be sent to the packet's original destination address instead of being redirected to the selected traffic handling system; and

sending the packet from the first computer system to its original destination.

44. (Original) A computer system as recited in claim 43, wherein it is determined that the packet is to be sent to the packet's original destination by determining that the packet is encapsulated and de-encapsulating the packet prior to sending the packet to its original destination.

45. (Previously presented) A computer system as recited in claim 40, wherein the traffic information sent from a selected traffic handling system includes service options specifying which data to be redirected to the selected traffic handling system.

46. (Original) A computer system as recited in claim 45, wherein the service options include a plurality of fields that are configurable to indicate that one or more fields of a packet received in the first computer system are to be used to determined redirection of packets to the selected traffic handling system.

47. (Original) A computer system as recited in claim 46, wherein the fields are selected from a group consisting of a source IP field, a destination IP field, a source port field, a destination port field, a source IP alternative field, a destination IP alternative field, a source port alternative field, and a destination port alternative field.

48. (Original) A computer system as recited in claim 47, wherein each field indicates that a corresponding field of a packet received in the first computer system is to be used to generate an index to a table identifying the plurality of associated traffic handling systems, the generated index being associated with the selected traffic handling system.

49. (Original) A computer system as recited in claim 48, wherein each field indicates that a hashed value of the corresponding field of the received packet is to be used to generate the index to the table identifying the plurality of associated traffic handling systems.

50. (Original) A computer system as recited in claim 46, wherein at least one of the fields may be set to indicate one or more port identifiers of traffic received in the computer system.

51. (Original) A computer system as recited in claim 50, wherein the fields are selected from a group consisting of a port 0 field, a port 1 field, a port 2 field, a port 3 field, a port 4 field, a port 5 field, a port 6 field, and a port 7 field.

52. (Original) A computer system as recited in claim 51, wherein the fields includes a source/destination field to indicate whether the port identifiers of the received traffic are source ports or destination ports.

53. (Original) A computer system as recited in claim 50, wherein at least one of the memory and the processor are further adapted to provide:

receiving a packet from a first processing device destined for a second processing device; and

when one or more port identifiers of the received packet matches a corresponding set field of the service options of the selected traffic handling system, redirecting the packet to the selected traffic handling system.

54. (Previously presented) A computer system as recited in claim 53, wherein the traffic information sent from a selected traffic handling system to the first computer system includes security options for specifying an authentication level for messages communicated between the first computer system and the selected traffic handling system.

55. (Original) A computer system as recited in claim 54, wherein the security options are configurable to select no authentication for messages communicated between the first computer system and the selected traffic handling system.

56. (Original) A computer system as recited in claim 54, wherein the security options are configurable to require a predetermined password encoded within messages communicated between the first computer system and the selected traffic handling system.

57. (Original) A computer system as recited in claim 40, wherein the traffic information sent from a selected traffic handling system includes identifying information for the selected traffic handling system.

58-77. (Cancelled)

78. (Previously presented) A computer readable storage medium, that is in the form of magnetic media, optical media, or magneto-optical media, for facilitating redirection of data sent

from a first processing device to a second processing device, comprising computer readable code that is stored on such computer readable storage medium and that is configured for performing the following operations:

receiving traffic information, at a third processing device associated with a plurality of traffic handling systems, from each of the associated traffic handling systems, wherein the traffic information received from each associated traffic handling system specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system;

communicating the traffic information to at least a designated one of the associated traffic handling systems;

receiving traffic redirection information, at the third processing device, from the designated traffic handling system, the traffic redirection information specifying which data is to be redirected to which one of the plurality of traffic handling systems and being based on the communicated traffic information;

receiving data, at the third processing device, from one or more first processing devices that are destined for one or more second processing devices; and

redirecting, at the third processing device, the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data.

79. (Cancelled)

80. (Previously presented) An apparatus for facilitating redirection of data sent from a first processing device to a second processing device, wherein the apparatus is associated with a plurality of traffic handling systems, the apparatus comprising:

means for receiving traffic information from each of the associated traffic handling systems, wherein the traffic information received from each of the associated traffic handling systems specifies which data based on at least a portion of the data should be redirected to the each associated traffic handling system;

means for determining how to redirect data received by the apparatus to a selected traffic handling system based on the received traffic information;

means for receiving data, at the third processing device, from one or more first processing devices that are destined for one or more second processing devices; and

means for redirecting, at the third processing device, the received data to selected one or more of the traffic handling systems so that the redirected data are apportioned between the traffic handling systems based on the traffic information from each of the associated traffic handling systems and at least a portion of the received data.

81-82. (Cancelled)

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

A copy of the Notice of Panel Decision from Pre-Appeal Brief Review was mailed on June 25, 2007 is appended herein.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/608,802	06/30/2000	Martin Cieslak	CISCP146	2424

22434 7590 06/25/2007

BEYER WEAVER LLP
P.O. BOX 70250
OAKLAND, CA 94612-0250

EXAMINER	
BHATIA, AJAY M	

ART UNIT	PAPER NUMBER
2145	

MAIL DATE	DELIVERY MODE
06/25/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Application Number	Application/Control No.	Applicant(s)/Patent under Reexamination	
	09/608,802	CIESLAK ET AL.	
	Jason D. Cardone	Art Unit 2145	

Document Code - AP.PRE.DEC

Notice of Panel Decision from Pre-Appeal Brief Review

This is in response to the Pre-Appeal Brief Request for Review filed 4/30/07.

1. ☐ **Improper Request** – The Request is improper and a conference will not be held for the following reason(s):

☐ The Notice of Appeal has not been filed concurrent with the Pre-Appeal Brief Request.
☐ The request does not include reasons why a review is appropriate.
☐ A proposed amendment is included with the Pre-Appeal Brief request.
☐ Other:

The time period for filing a response continues to run from the receipt date of the Notice of Appeal or from the mail date of the last Office communication, if no Notice of Appeal has been received.

2. ☒ **Proceed to Board of Patent Appeals and Interferences** – A Pre-Appeal Brief conference has been held. The application remains under appeal because there is at least one actual issue for appeal. Applicant is required to submit an appeal brief in accordance with 37 CFR 41.37. The time period for filing an appeal brief will be reset to be one month from mailing this decision, or the balance of the two-month time period running from the receipt of the notice of appeal, whichever is greater. Further, the time period for filing of the appeal brief is extendible under 37 CFR 1.136 based upon the mail date of this decision or the receipt date of the notice of appeal, as applicable.

☒ The panel has determined the status of the claim(s) is as follows:
Claim(s) allowed: *None*.
Claim(s) objected to: *None*.
Claim(s) rejected: *1-19, 40-57, 78 and 80*.
Claim(s) withdrawn from consideration: *None*.

3. ☐ **Allowable application** – A conference has been held. The rejection is withdrawn and a Notice of Allowance will be mailed. Prosecution on the merits remains closed. No further action is required by applicant at this time.

4. ☐ **Reopen Prosecution** – A conference has been held. The rejection is withdrawn and a new Office action will be mailed. No further action is required by applicant at this time.

JASON CARDONE
SUPERVISORY PATENT EXAMINER

All participants:

(1) *Jason D. Cardone*.
(2) *Ajay Bhatia*.
(3) *William Vaughn*.
(4)_____.